**Arxiv manuscript No.**
(will be inserted by the editor)

# Collecting and Annotating the Large Continuous Action Dataset

**Daniel Paul Barrett** · **Ran Xu** · **Haonan Yu** · **Jeffrey Mark Siskind**

arXiv:1511.05914v1 [cs.CV] 18 Nov 2015

Received: date / Accepted: date

**Abstract** We make available to the community a new dataset to support action-recognition research. This dataset is different from prior datasets in several key ways. It is significantly larger. It contains streaming video with long segments containing multiple action occurrences that often overlap in space and/or time. All actions were filmed in the same collection of backgrounds so that background gives little clue as to action class. We had five humans replicate the annotation of temporal extent of action occurrences labeled with their class and measured a surprisingly low level of intercoder agreement. A baseline experiment shows that recent state-of-the-art methods perform poorly on this dataset. This suggests that this will be a challenging dataset to foster advances in action-recognition research. This manuscript serves to describe the novel content and characteristics of the LCA dataset, present the design decisions made when filming the dataset, and document the novel methods employed to annotate the dataset.

**Keywords** action recognition · dataset · video

**Mathematics Subject Classification (2000)** 68T45

Daniel Paul Barrett
School of Electrical and Computer Engineering, Purdue University

Ran Xu
Computer Science and Engineering, SUNY Buffalo

Haonan Yu
School of Electrical and Computer Engineering, Purdue University

Jeffrey Mark Siskind
School of Electrical and Computer Engineering, Purdue University
Tel.: +1-765-496-3197
Fax: +1-765-494-2706
E-mail: qobi@purdue.edu

## 1 Introduction

There has been considerable research interest in action recognition in video over the past two decades [2–5, 7, 8, 11–19, 21–28, 31, 34–36, 38–40, 42–46, 48–52, 54, 56, 57, 59–71, 73]. To support such research, numerous video datasets have been gathered. Liu *et al.* [33] summarize the available datasets as of 2011. These include KTH (6 classes, [49]), Weizmann (10 classes, [4]), CMU Soccer (7 classes, [9]), CMU Crowded (5 classes, [21]), UCF Sports (9 classes, [45]), UR ADL (10 classes, [38]), UM Gesture (14 classes, [32]), UCF Youtube (11 classes, [35]), Hollywood-1 (8 classes, [29]), Hollywood-2 (12 classes, [37]), MultiKTH (6 classes, [58]), MSR (3 classes, [72]), and TRECVID (10 classes, [53]). These datasets contain short clips, each depicting one of a small number of classes (3–14). Several more recent datasets also contain short clips, each depicting a single action, but with a larger number of action classes: UCF50 (50 classes, [44]), HMDB51 (51 classes, [27]), and UCF101 (101 classes, [55]). The VIRAT dataset [41] has 12 classes and longer streaming video.

Here, we introduce a new dataset called the *Large Continuous Action Dataset* (LCA). This dataset contains depictions of 24 action classes. The video for this dataset was filmed and annotated as part of the DARPA Mind's Eye program. A novel characteristic of this dataset is that rather than consisting of short clips each of which depicts a single action class, this dataset contains much longer streaming video segments that each contain numerous instances of a variety of action classes that often overlap in time and may occur in different portions of the field of view. The annotation that accompanies this dataset delineates not only which actions occur but also their temporal extent.

Many of the prior datasets were culled from video downloaded from the internet. In contrast, the LCA dataset contains video that was filmed specifically to construct the dataset.

While the video was filmed with people hired to act out the specified actions according to a general script, the fact that the video contains long streaming segments tends to mitigate any artificial aspects of the video and render the action depictions to be quite natural. Moreover, the fact that all of the video was filmed in a relatively small number of distinct backgrounds makes the dataset challenging; the background gives little clue as to the action class.

A further distinguishing characteristic of the LCA dataset is the degree of ambiguity. Most prior action-recognition datasets, in fact most prior datasets for all computer-vision tasks, make a tacit assumption that the labeling is unambiguous and thus there is a 'ground truth.' We had a team of five human annotators each annotate the entire LCA dataset. This allowed us to measure the degree of intercoder agreement. Surprising, there is a significant level of disagreement between humans as to the temporal extent of most action instances. We believe that such inherent ambiguity is a more accurate reflection of the underlying action-recognition task and hope that the multiplicity of divergent annotations will help spur novel research with this more realistic dataset.

Another distinguishing characteristic of the LCA dataset is that some action occurrences were filmed simultaneously with multiple cameras with partially overlapping fields of view. While the cameras were neither spatially calibrated nor temporally synchronized, the fact that we have multiple annotations of the temporal extent of action occurrences may support future efforts to perform temporal synchronization after the fact. Furthermore, while most of the video was filmed from ground-level cameras with horizontal view, some of the video was filmed with aerial cameras with bird's eye view. Some of this video was filmed simultaneously with ground cameras. This may support future efforts to conduct scene reconstruction.

Some datasets are provided with specific tasks and evaluation metrics. We refrain from doing so for this dataset. *Inter alia*, we do not provide official sanctioned splits into validation sets. Instead, we leave it up to the community to make use of this dataset in a creative fashion for as many different tasks as it will be suited.

In particular, the evaluations conducted by Mind's Eye included a specific set of tasks, namely Recognition (REC), Description (DES), Gap Filling (GAP), and Anomaly Detection (ANM) with specific evaluation metrics. Such tasks and metrics are expressly *not* part of LCA. The evaluations conducted under Mind's Eye make use of material that is not included in LCA and metrics that are not public. Likewise, LCA contains material that was not available for use during the Mind's Eye evaluations. Thus said evaluations could not be replicated outside of the context of Mind's Eye. Any potential future evaluations conducted with LCA would thus be incomparable to the results obtained under Mind's Eye.

The entire LCA dataset, including the video and the annotations, has been cleared for release by DARPA. The remaining material gathered by DARPA for the Mind's Eye Year 2 evaluation that is not included in LCA may not have been cleared for release. As part of the release process, some video was edited to remove certain portions. Furthermore, the annotation process was performed with the particular versions of the videos included in LCA as provided by DARPA. These may have been transcoded from the original as filmed by the camera. Thus, the time alignment of the annotations can only be guaranteed with the versions of the videos included in LCA. The time alignment may not be correct for any other versions of these videos that may be residual from the DARPA Mind's Eye program.

## 2 Collection

The video for this dataset was filmed by DARPA in conjunction with Mitre and several performers from the Mind's Eye program.[1] The bulk of the video was filmed as part of the Mind's Eye Year 2 evaluation. Within the Mind's Eye program, that video was referred to as C-D2a, C-D2b, C-D2c, and the Y2 Evaluation dataset. This video is disjoint from that gathered by Janus Research Group as part of the Mind's Eye Year 1 evaluation. Within the Mind's Eye program, that video was refereed to as C-D1a, C-D1b, C-D1, and C-E1. As the LCA dataset contains only a subset of that material, we refrain from using all such terminology in reference to LCA. Also note that the LCA dataset contains some video that was not included in the data used as part of the Mind's Eye evaluations.

The LCA dataset was filmed at three different locations over four periods:

1. The Great Plains Joint Training Center (GPJTC), an army training facility in Kansas, on 22 August 2011. Filming took place in two contexts at GPJTC, a simulated country road and a simulated safe house.
2. Strategic Operations, Inc. (STOPS), a training facility in San Diego, on 14–15 December 2011 and on 6–9 March 2012. Filming during the first period took place in three contexts: two different simulated country roads and one simulated safe house. Filming during the second period took place in five contexts: two different simulated country roads, two different simulated safe houses, and one other.
3. Fort Indiantown Gap (FITG), an army training facility in Pennsylvania, on 19–20 June 2012.

A portion of the video was annotated by Mitre for the Mind's Eye Year 2 evaluation. That annotation is not included in LCA. After the completion of the Mind's Eye

[1] `http://www.visint.org/`

**Table 1** Verbs used as labels in the LCA dataset. The starred verbs were used as part of the stage directions to the actors. The remaining verbs were not used as part of the stage directions but may have occurred incidentally.

| | | | |
|---|---|---|---|
| *approach** | *drop** | *give** | *replace** |
| *arrive* | *enter** | *hold* | *run* |
| *bury** | *exchange** | *leave* | *stop* |
| *carry** | *exit** | *pass** | *take** |
| *chase** | *flee** | *pick up** | *turn* |
| *dig** | *follow** | *put down** | *walk* |

Year 2 evaluation, we undertook a systematic annotation effort for a portion of the above video. That annotation forms the basis of LCA. LCA contains all and only the portion of the above video that was annotated as part of this process. This video comprises 190 files as delineated in Table 2. Eight files are MOV format, 46 are MP4 format, and 136 are AVI format. The MOV files all use the MP4V codec and are 640×384 at 60 fps. The MP4 files all use the H264 codec and are 640×360 at 30 fps. The AVI files all use the XVID codec and are either 640×384 at 60 fps or 1440×1080 at 30 fps. This constitutes 2302144 frames and a total of 12 hours, 51 minutes, and 16 seconds of video. For comparison, UCF50 has 1330936 frames and 13.81 hours, HMDB51 has 632635 frames and 5.85 hours, UCF101 has 27 hours, Hollywood-2 has 20.1 hours, and VIRAT has 8.5 hours. Several frame sequences from this dataset illustrating several of the backgrounds are shown in Fig. 1.

The Mind's Eye program specified a set of 48 verbs of interest. Of these, the LCA dataset uses only 24 verbs as annotation labels, as delineated in Table 1. Of these, 17 verbs were used as part of the stage directions given to the actors to guide the actions that they performed. The remainder were not used as part of the stage directions but occurred incidentally. Nothing, however, precluded the actors from performing actions that could be described by other verbs. Thus the video depicts many other actions than those annotated, including but not limited to riding bicycles, pushing carts, singing, pointing guns, arguing, and kicking balls. The only restriction, in principle, to these 24 verbs is that these were the only actions that were annotated. Identifying the presence of specific verbs in the context of many such confounding actions should present additional challenges.

## 3 Annotation

We annotated all occurrences of the 24 verbs from Table 1 in the videos in Table 2. Each such occurrence consisted of a temporal interval labeled with a verb. The judgment of whether an action described by a particular verb occurred is subjective; different annotators will arrive at different judgments as to occurrence as well as the temporal extent thereof. To help guide annotators, we gave them the specification of the intended meaning of each of the 24 verbs as provided by DARPA. Annotators performed the annotation at workstations with dual monitors. One monitor displayed the annotation tool while the other monitor displayed the documentation of intended verb meaning. The documentation of intended verb meaning is included in the LCA distribution.

We also asked annotators to annotate intervals where certain object classes were present in the field of view. These include *bandannas*, *bicycles*, *people*, *vehicles*, and *weapons*. (The *bandannas* were worn by *people* around their head or arms.) For these, a count of the number of instances of each class that were visible in the field of view was maintained. It was incremented each time a new instance became visible and decremented each time an instance became invisible. We instructed annotators that there was no need to be precise when an instance was partially visible. We further instructed annotators that *vehicles* denoted motor vehicles, not push carts, and *weapons* denoted guns, not other things like clubs or rocks that could be used as weapons.

We provided annotators with a tool that allowed them to view the videos at ordinary frame rate, stop and start the videos at will, navigate to arbitrary points in the videos, view individual frames of the videos, add, delete, and move starting and ending points of intervals, and label intervals with verbs. The tool also contained buttons to increment and decrement the counts for each of the object classes and appraised the annotator with the running counts for the object classes in each frame as the video was played or navigated.

Because of the large quantity of video to be annotated, and the fact that nothing happens during large portions of the video, we preprocessed the video to reduce the amount requiring manual annotation. We first downsampled the video to 5 fps just for the purpose of annotation; the annotation was converted back at the end to the original frame rate. Then segments of this downsampled video where no motion occurred were removed. To do this, we computed dense optical flow on each pixel of each frame of the downsampled video. We then computed the average of the magnitude of the flow vectors in each frame and determined which frames were above a threshold. Stretches of contiguous frames that were above threshold that were separated by short stretches of contiguous frame that were below threshold were merged into single temporal segments. Then such single temporal segments that were shorter than a specified temporal length were discarded.[2] Annotators were only given the remaining temporal segments to annotate. We performed a postprocessing step whereby the authors manually viewed all discarded frames to make sure that no actions started, ended, or spanned the omitted temporal segments. As part of this post

[2] The threshold for average optical flow magnitude was 150. The threshold for ignoring short below-threshold spans when merging contiguous above-threshold frames into temporal segments was 50 frames. The threshold for ignoring short temporal segments was 15 frames.

**Table 2** The original names of the files provided by DARPA. Filenames containing GPTC were filmed at GPJTC. Filenames containing STOPS were filmed at STOPS. Filenames consisting solely of a number were filmed at FITG. Numbers of the form YYYYMMDD indicate filming date. CR indicates country road. SH indicates safe house. Indices on CR, SH, and VT indicate variant backgrounds of the given class. CP1, CP2, C1, and C3 indicate camera. Text indicates the staging directions to guide filming. The remaining numbers serve to uniquely identify the video. These videos were renamed to `video-XXX` for consistency in the release. The release also contains a file, `video-mapping.txt`, which includes the mapping between the original filenames and those in the LCA release.

| | | | |
|---|---|---|---|
| GPTC_20110822_SH_02_CP1_NOACTIVITY | STOPS_20120307_SH3_06_C1-edited-01 | STOPS_20120308_CR1_06a_C1 | 100 |
| GPTC_20110822_SH_02_CP2_NOACTIVITY | STOPS_20120307_SH3_06_C3-edited-01 | STOPS_20120308_CR1_06a_C3 | 105 |
| GPTC_20110822_SH_07_CP1_EX&RET | STOPS_20120307_SH3_03_C1-edited-01 | STOPS_20120306_SH1_04_C1 | 114 |
| GPTC_20110822_SH_07_CP2_EX&RET | STOPS_20120307_SH3_03_C3-edited-01 | STOPS_20120306_SH1_04_C3 | 118 |
| GPTC_20110822_SH_13_CP1_WARYHO | STOPS_20120307_SH3_02_C1-edited-01 | STOPS_20120306_SH1_01_C1 | 11 |
| GPTC_20110822_SH_13_CP2_WARYHO | STOPS_20120307_SH3_02_C3-edited-01 | STOPS_20120306_SH1_01_C3 | 121 |
| GPTC_20110822_SH_11_CP1_PARAHO | STOPS_20120307_SH3_01_C1-edited-01 | STOPS_20120308_CR1_13a_C1 | 128 |
| GPTC_20110822_SH_11_CP2_PARAHO | STOPS_20120307_SH3_01_C3-edited-01 | STOPS_20120308_CR1_13a_C3 | 133 |
| GPTC_20110822_SH_12_CP1_HOSTILEHO | STOPS_20120308_CR1_07a_C1 | STOPS_20120309_VT1_07_C1 | 154 |
| GPTC_20110822_SH_12_CP2_HOSTILEHO | STOPS_20120308_CR1_07a_C3 | STOPS_20120309_VT1_07_C3 | 158 |
| GPTC_20110822_SH_06_CP1_SUPCACHRET | STOPS_20120309_VT1_23_C1 | STOPS_20120309_VT1_25_C1 | 162 |
| GPTC_20110822_SH_06_CP2_SUPCACHRET | STOPS_20120309_VT1_23_C3 | STOPS_20120309_VT1_25_C3 | 175 |
| GPTC_20110822_SH_09_CP1_GIVE&CONT-BLDG | STOPS_20120308_CR1_NA_C1 | STOPS_20120309_VT1_05_C1 | 188 |
| GPTC_20110822_SH_09_CP2_GIVE&CONT-BLDG | STOPS_20120308_CR1_NA_C3 | STOPS_20120309_VT1_05_C3 | 195 |
| GPTC_20110822_SH_05_CP1_SUPCACHDUMP | STOPS_20120306_SH1_06_C1 | STOPS_20120308_CR1_01c_C1 | 196 |
| GPTC_20110822_SH_05_CP2_SUPCACHDUMP | STOPS_20120306_SH1_06_C3 | STOPS_20120308_CR1_01c_C3 | 200 |
| GPTC_20110822_SH_08_CP1_HO&EX-mixBKG | STOPS_20120306_SH1_NA_C1 | STOPS_20120309_VT1_NA_C1 | 211 |
| GPTC_20110822_SH_08_CP2_HO&EX-mixBKG | STOPS_20120306_SH1_NA_C3 | STOPS_20120309_VT1_NA_C3 | 226 |
| GPTC_20110822_SH_03_CP1_HO&RET | STOPS_20120308_CR1_08a_C1 | STOPS_20120308_CR1_03a_C1 | 22 |
| GPTC_20110822_SH_03_CP2_HO&RET | STOPS_20120308_CR1_08a_C3 | STOPS_20120308_CR1_03a_C3 | 230 |
| GPTC_20110822_SH_04_CP1_HO&RET2 | STOPS_20120308_CR1_12a_C1 | STOPS_20120306_SH1_02_C1 | 238 |
| GPTC_20110822_SH_04_CP2_HO&RET2 | STOPS_20120308_CR1_12a_C3 | STOPS_20120306_SH1_02_C3 | 243 |
| GPTC_20110822_CR_13_CP1_NOACTIVITY | STOPS_20120308_CR1_01a_C1 | STOPS_20120307_SH3_11_C1 | 252 |
| GPTC_20110822_CR_13_CP2_NOACTIVITY | STOPS_20120308_CR1_01a_C3 | STOPS_20120307_SH3_11_C3 | 265 |
| GPTC_20110822_CR_07_CP1_SALESMAN | STOPS_20120309_VT1_01_C1 | STOPS_20120308_CR1_04a_C1 | 269 |
| GPTC_20110822_CR_07_CP2_SALESMAN | STOPS_20120309_VT1_01_C3 | STOPS_20120308_CR1_04a_C3 | 277 |
| GPTC_20110822_CR_12_CP1_BAGDOWNHO&RET | STOPS_20120309_VT1_26_C1 | STOPS_20120307_SH3_10_C1 | 30 |
| GPTC_20110822_CR_12_CP2_BAGDOWNHO&RET | STOPS_20120309_VT1_26_C3 | STOPS_20120307_SH3_10_C3 | 39 |
| GPTC_20110822_CR_02_CP1_RoutineActivity | STOPS_20120308_CR1_02a_C1 | STOPS_20120309_VT1_22_C1 | 41 |
| GPTC_20110822_CR_02_CP2_RoutineActivity | STOPS_20120308_CR1_02a_C3 | STOPS_20120309_VT1_22_C3 | 46 |
| GPTC_20110822_CR_05_CP1_CONTHUR-NOHO | STOPS_20120308_CR1_05a_C1 | STOPS_20120308_CR1_01b_C1 | 52 |
| GPTC_20110822_CR_05_CP2_CONTHUR-NOHO | STOPS_20120308_CR1_05a_C3 | STOPS_20120308_CR1_01b_C3 | 57 |
| GPTC_20110822_CR_04_CP1_EX | STOPS_20120309_VT1_04_C1 | STOPS_20120309_VT1_28_C1 | 59 |
| GPTC_20110822_CR_04_CP2_EX | STOPS_20120309_VT1_04_C3 | STOPS_20120309_VT1_28_C3 | 62 |
| GPTC_20110822_CR_08_CP1_SALESMAN_INSIST | STOPS_20120306_SH1_07_C1 | STOPS_20120309_VT1_20_C1 | 68 |
| GPTC_20110822_CR_08_CP2_SALESMAN_INSIST | STOPS_20120306_SH1_07_C3 | STOPS_20120309_VT1_20_C3 | 69 |
| GPTC_20110822_CR_11_CP1_DOUBLEAVOIDBADGUY | STOPS_20120309_VT1_21_C1 | STOPS_20120309_VT1_03_C1 | 81 |
| GPTC_20110822_CR_11_CP2_DOUBLEAVOIDBADGUY | STOPS_20120309_VT1_21_C3 | STOPS_20120309_VT1_03_C3 | 86 |
| GPTC_20110822_CR_09_CP1_UPTONOGOOD | STOPS_20120307_SH3_07_C1 | STOPS_20120309_VT1_27_C1 | 88 |
| GPTC_20110822_CR_09_CP2_UPTONOGOOD | STOPS_20120307_SH3_07_C3 | STOPS_20120309_VT1_27_C3 | 97 |
| GPTC_20110822_CR_06_CP1_DISAGREE_NOGIVE_RETURN | STOPS_20120308_CR1_09a_C1 | STOPS_20120309_VT1_13_C1 | |
| GPTC_20110822_CR_06_CP2_DISAGREE_NOGIVE_RETURN | STOPS_20120308_CR1_09a_C3 | STOPS_20120309_VT1_13_C3 | |
| GPTC_20110822_CR_10_CP1_AVOIDBADGUY | STOPS_20120307_SH3_08_C1 | STOPS_20120308_CR1_06b_C1 | |
| GPTC_20110822_CR_10_CP2_AVOIDBADGUY | STOPS_20120307_SH3_08_C3 | STOPS_20120308_CR1_06b_C3 | |
| GPTC_20110822_CR_03_CP1_HO | STOPS_20120308_CR1_11a_C1 | STOPS_20120309_VT1_02_C1 | |
| GPTC_20110822_CR_03_CP2_HO | STOPS_20120308_CR1_11a_C3 | STOPS_20120309_VT1_02_C3 | |
| | STOPS_20120306_SH1_05_C1 | STOPS_20120308_CR1_10a_C1 | |
| | STOPS_20120306_SH1_05_C3 | STOPS_20120308_CR1_10a_C3 | |
| | STOPS_20120309_VT1_12_C1 | STOPS_20120309_VT1_24_C1 | |
| | STOPS_20120309_VT1_12_C3 | STOPS_20120309_VT1_24_C3 | |
| | STOPS_20120309_VT1_10_C1 | STOPS_20120308_CR1_03b_C1 | |
| | STOPS_20120309_VT1_10_C3 | STOPS_20120308_CR1_03b_C3 | |

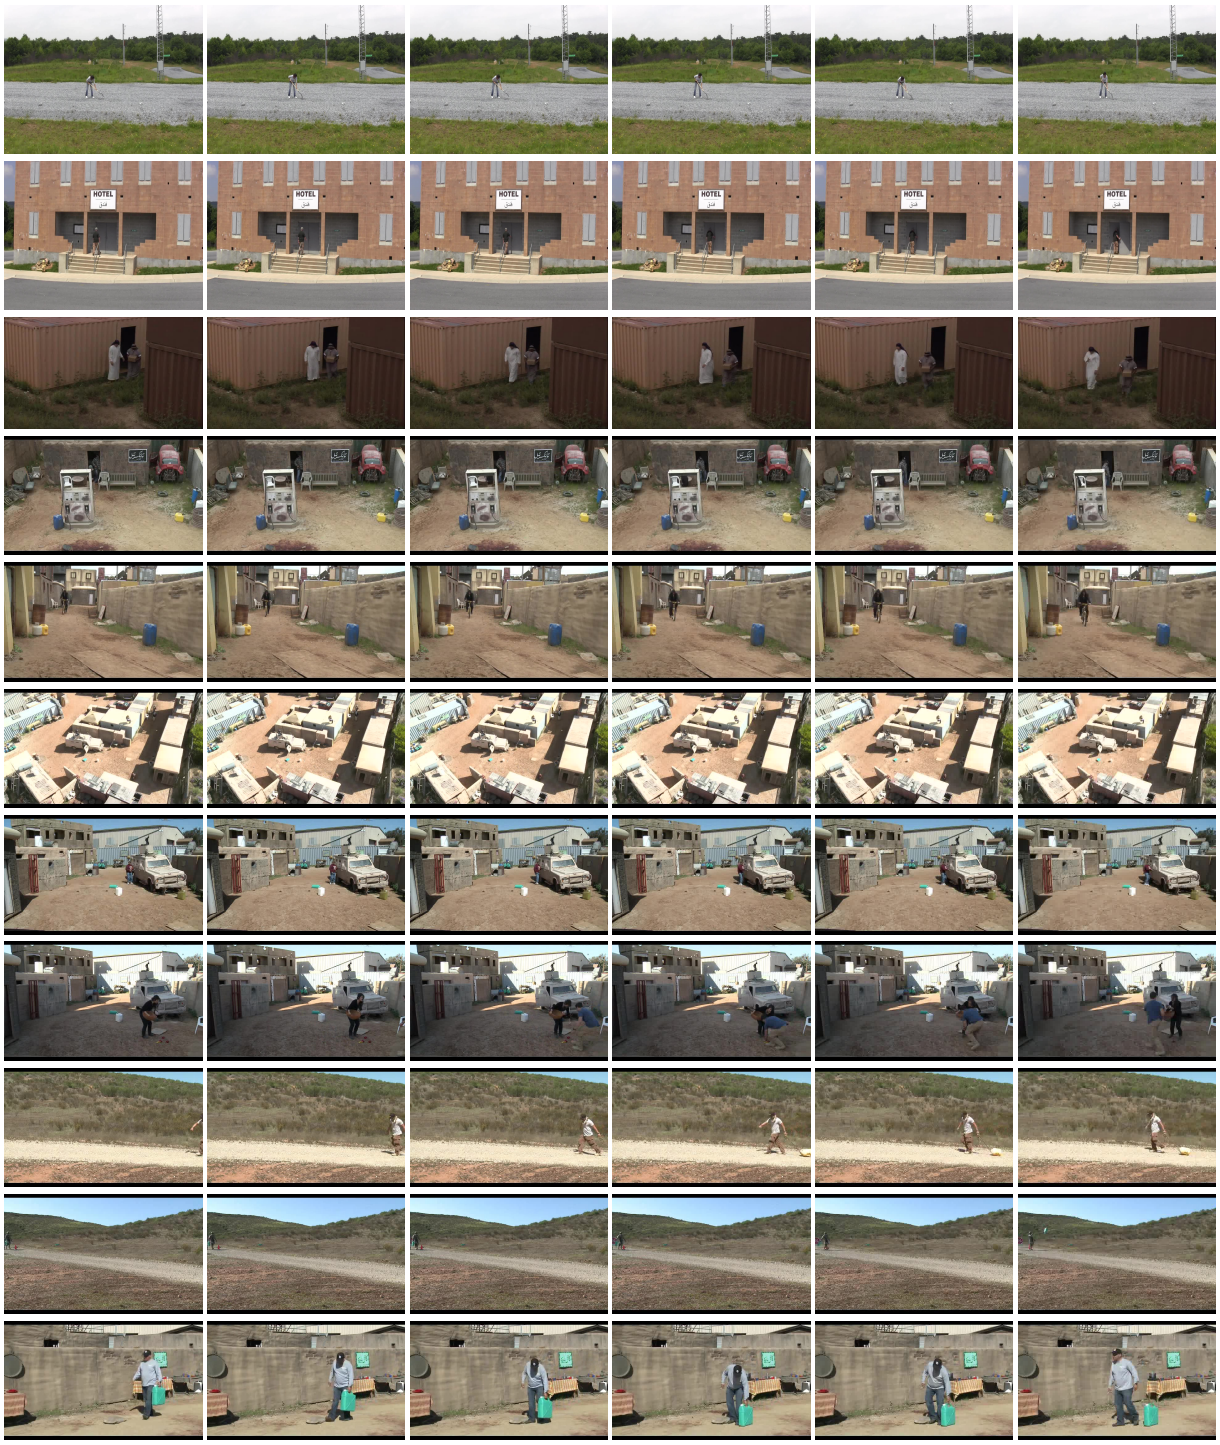

**Fig. 1** Several frame sequences from the LCA dataset illustrating several of the backgrounds in which they were filmed.

processing step, the authors manually checked that none of the specified object classes entered or left the field of view during the omitted temporal segments.

We had five annotators each independently annotate the entire LCA dataset. Annotators were given initial instructions. During the annotation, annotators were encouraged to discuss their annotation judgments with the authors. The authors would then arbitrate the judgment, often specifying principles to guide the annotation. These principles were then circulated among the other annotators. The annotator instructions and principles developed through arbitration are included in the LCA distribution.

We performed a consistency check during the annotation process. Whenever an annotator completed annotation

of a temporal segment, if that annotator did not annotate any intervals during that segment but other annotators did, we asked that annotator to review their annotation.

The LCA dataset contains five verb-annotation files for each of the video files in Table 2. These have the same name as their corresponding video, but with the extension `txt`, and are located in directories named with each of the annotator codes `bmedikon`, `cbushman`, `kim861`, `nielder`, and `nzabikh`. Each line in each of these files contains a single temporal interval as a text string specifying a verb and two zero-origin nonnegative integers specifying the starting and ending frames of the interval inclusive. The LCA dataset also contains five object-class annotation files for each of the video files in Table 2. These also share the filename with the corresponding video, but with the addition of the suffix `-enter-exits.txt`, and are located in the same directories named with each of the above annotator codes. Each line in each of these files contains a text string specifying an object class, an integer specifying the number of objects of that class entering or exiting the field of view (positive for entering and negative for exiting), and a single zero-origin nonnegative integer specifying the video frame.

## 4 Analysis

We analyzed the degree of agreement between the different annotators. To do this, we compared pairs of annotators, taking the judgments of one as 'ground truth' and computing the F1 score of the other. An interval in the annotation being scored was taken as a true positive if it overlapped some interval with the same label in the 'ground truth.' An interval in the annotation being scored was taken as a false positive if it didn't overlap any interval with the same label in the 'ground truth.' An interval in the 'ground truth' was taken as a false negative if it didn't overlap any interval with the same label in the annotation being scored. From these counts, an F1 score could be computed.

We employed the following overlap criterion. For a pair of intervals, we computed a one-dimensional variant of the 'intersection over union' criterion employed within the Pascal VOC Challenge to determine overlap of two axis-aligned rectangles [10], namely the length of the intersection divided by the length of the union. We considered two intervals to overlap when the above exceeded some specified threshold. We then computed the F1 score as this threshold was varied and plotted the results for all pairs of annotators (Fig. 2).

Note that there is a surprisingly low level of agreement between annotators. Annotators rarely if ever agree on the precise temporal extent of an action as indicated by the fact that all agreement curves go to zero as the overlap threshold goes to one. At an overlap threshold of 0.5, the F1 score varies between about 0.3 and about 0.6. At an overlap threshold of 0.1, the threshold employed by VIRAT to score machines against humans, the F1 score varies between about 0.38 and about 0.67. This would put an upper bound on machine performance with this dataset using the VIRAT threshold. Even if the overlap threshold is reduced to zero, the F1 score varies between about 0.43 and about 0.7. This indicates that this dataset should be challenging for computer action recognition.

This difficulty is corroborated by a recent paper [1]. That paper employs a different subset of video from the DARPA Mind's Eye Year 2 evaluation that is extremely similar to that in the LCA dataset. That dataset was annotated with the same procedures that were used to annotate the LCA dataset. The six verbs with highest intercoder agreement were selected: *carry*, *dig*, *hold*, *pick up*, *put down*, and *walk*. For each of these, between 23 and 30 clips of 2.5s duration with the highest level of intercoder agreement were selected, yielding 169 distinct clips. Seven different state-of-the-art computer-vision action-recognition methods (C2 [20], Action Bank [48], Stacked ISA [30], VHTK [38], Cao's implementation [6] of Ryoo's method [47], Cao's method [6], and our own implementation of the classifier described by Wang *et al.* [61] on top of the Dense Trajectories [60–62] feature extractor) were employed on this dataset, performing one-of-out-six classification in an eight-fold cross-validation. Note that for this task, each 2.5s clip was labeled with precisely one of the six verbs as ground truth. All seven methods performed extremely poorly on this dataset (C2 47.4%, Action Bank 44.2%, Stacked ISA 46.8%, VHTK 32.5%, Cao's implementation of Ryoo's method 31.2%, Cao's method 33.3%, and Dense Trajectories 52.3%), a task with only six classes and chance performance of 16.6%.

## 5 Baseline Experiment

We performed an experiment similar to that of Barbu *et al.* [1] to present and compare the performance of several state-of-the-art action-recognition systems on the LCA dataset. We evaluated all known action-recognition systems for which the code for the end-to-end system is available, as well as implementations of several for which the code is unavailable. We used the same set of seven state-of-the-art action-recognition systems compared in Barbu *et al.* [1] (C2 [20], Action Bank [48], Stacked ISA [30], VHTK [38], Cao's implementation [6] of Ryoo's method [47], Cao's method [6], and our own implementation of the classifier described by Wang *et al.* [61] on top of the Dense Trajectories [60, 61] feature extractor) We also compared against our own implementation of the classifier described by Wang *et al.* [61] on top of the more recent Improved Trajectories method [62].

As these methods are designed for classification of video clips, rather than for streaming video, this experiment was performed on a subset of the LCA dataset. This subset was designed to be similar in character to other action-recognition

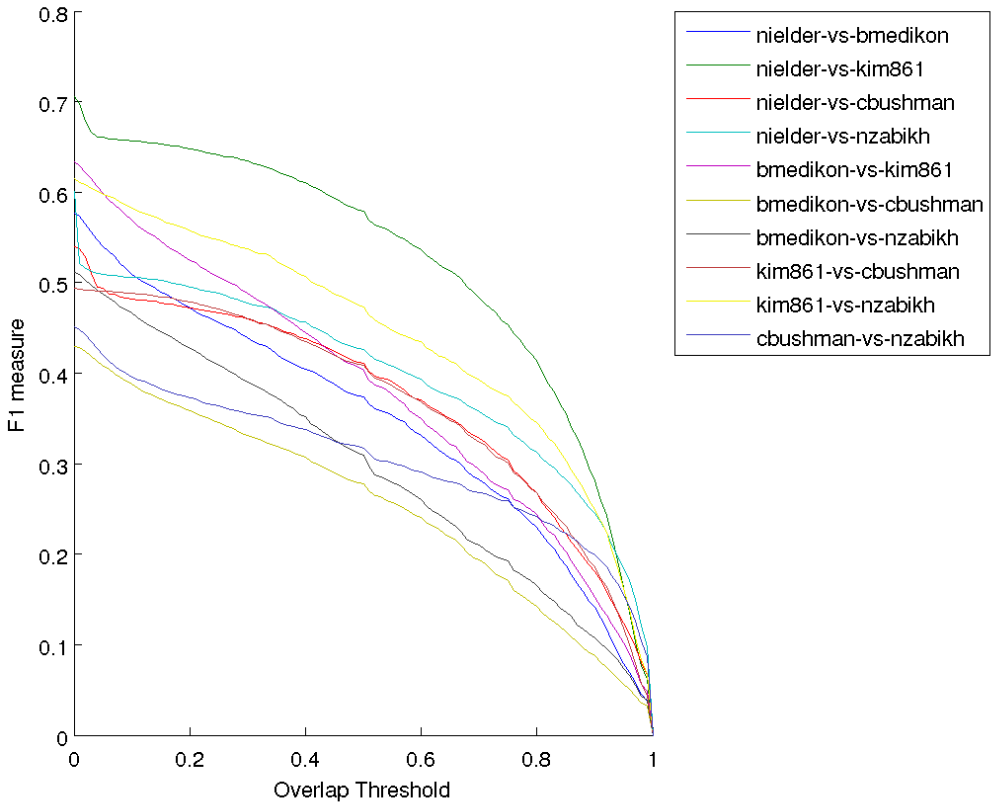


**Fig. 2** Intercoder agreement on the annotations of the LCA dataset. F1 score for each pair of annotators as the overlap criterion is varied. Overlap of two intervals is measured as the length of their intersection divided by the length of their union.

datasets and comprised short video clips. It was created as follows. First, we took the human-annotated action intervals produced by one of the annotators, `cbushman`. This annotator was chosen to maximize the number of available action intervals. Next, a maximum of 100 intervals were selected for each action class. For those action classes for which more than 100 intervals were annotated, a random subset of 100 intervals was selected. For those action classes with 100 or fewer annotated intervals, all annotated intervals were used. A 2s clip was extracted from the original videos centered in time on the middle of each selected annotation interval. These clips were temporally downsampled to 20 fps and spatially downsampled to a width of 320 pixels, maintaining the aspect ratio. This process resulted in a total of 1858 clips used for the baseline experiment.

The class label of each clip was considered to be the action class corresponding to the human-annotated interval from which the clip was derived. The clips for each class were randomly split into a training set with 70% of the clips and a test set with 30% of the clips, under the constraint that sets of clips extracted from the same video should fall completely into either the training or test set. This was done to avoid having clips from the same action (*e.g.*, two clips from the same person digging in the same location) from appearing in both the training and test sets. This resulted in a training set of 1318 training clips and 540 test clips. Each method was trained on the training set and used to produce labels on the test set. All methods were run with default or recommended parameters. These labels were compared to the intended class labels to measure the accuracy of each method. The results of this experiment are summarized in Table 3.

**Table 3** Comparison of accuracy for state-of-the-art action-recognition systems on a subset of the LCA dataset.

| Method | accuracy (%) |
|---|---|
| Action Bank [48] | 16.667 |
| Improved Trajectories [62] | 15.556 |
| Dense Trajectories [60, 61] | 14.074 |
| C2 [20] | 9.259 |
| Cao [6] | 7.592 |
| Cao's [6] implementation of Ryoo [47] | 6.667 |
| Stacked ISA [30] | 6.667 |
| VHTK [38] | 6.296 |
| chance (30/540) | 5.555 |

There are several things of note in these results. First, all the accuracies are quite low, indicating the difficulty of the LCA dataset. The highest performing method, Action Bank, is correct only 16.667% of the time. The four lowest performing methods have accuracies approaching chance performance (5.555%). Additionally, the newer methods do not necessarily outperform the older methods. C2 significantly outperforms four more recently published methods, while Action Bank is the best, outperforming even Improved Trajectories, which has the highest performance on several well known datasets including HMDB (57.2% vs Action Bank's 26.9%) and UCF50 (91.2%). We suspect that this difference in relative performance compared to other datasets is the result of the lack of correlation between background and action class which is often present in other datasets, as well as the presence of multiple people in the field of view and the small relative size of the people in the field of view. That the performance is so low and that the highest scoring methods on other datasets are not necessarily the same here shows that this dataset presents new and difficult challenges not present in other datasets.

## 6 Conclusion

Upon acceptance of this manuscript we will make available to the community a new dataset to support action-recognition research. This dataset has more hours of video than HMDB51, roughly the same amount of video as UCF50, about half as much video as UCF101 and Hollywood-2, but unlike these has streaming video and has about twice as much video and twice as many classes as VIRAT, the largest dataset of streaming video. A distinguishing characteristic of this dataset is that the video is streaming; long video segments contain many actions that start and stop at arbitrary times, often overlapping in space and/or time. A further distinguishing characteristic is that while all actions were filmed in a variety of backgrounds, every action occurs in every background so that background gives little information as to action class. We employed novel techniques to annotate the temporal extent of action occurrences. A multiplicity of human annotations allows measuring intercoder agreement. The above characteristics together with the surprisingly low level of intercoder agreement suggest that this will be a challenging dataset. This is confirmed by the low performance of recent methods on a baseline experiment which also shows that those methods which perform best on other datasets do not necessarily outperform other methods on this dataset. The new difficulties posed by this dataset should spur significant advances in action-recognition research.

**Acknowledgements** This research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-10-2-0060. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either express or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes, notwithstanding any copyright notation herein.

## References

1. Barbu, A., Barrett, D.P., Chen, W., Siddharth, N., Xiong, C., Corso, J.J., Fellbaum, C.D., Hanson, C., Hanson, S.J., Hélie, S., Malaia, E., Pearlmutter, B.A., Siskind, J.M., Talavage, T.M., Wilbur, R.B.: Seeing is worse than believing: Reading people's minds better than computer-vision methods recognize actions. In: European Conference on Computer Vision, pp. 612–627 (2014)
2. Barbu, A., Bridge, A., Burchill, Z., Coroian, D., Dickinson, S., Fidler, S., Michaux, A., Mussman, S., Siddharth, N., Salvi, D., Schmidt, L., Shangguan, J., Siskind, J.M., Waggoner, J., Wang, S., Wei, J., Yin, Y., Zhang, Z.: Video in sentences out. In: Uncertainty in Artificial Intelligence, pp. 102–112 (2012)
3. Barbu, A., Siddharth, N., Michaux, A., Siskind, J.M.: Simultaneous object detection, tracking, and event recognition. Advances in Cognitive Systems **2**, 203–220 (2012)
4. Blank, M., Gorelick, L., Shechtman, E., Irani, M., Basri, R.: Actions as space-time shapes. In: International Conference on Computer Vision, vol. 2, pp. 1395–1402 (2005)
5. Brand, M., Oliver, N., Pentland, A.: Coupled hidden Markov models for complex action recognition. In: Computer Vision and Pattern Recognition, pp. 994–999 (1997)
6. Cao, Y., Barrett, D., Barbu, A., Narayanaswamy, S., Yu, H., Michaux, A., Lin, Y., Dickinson, S., Siskind, J.M., Wang, S.: Recognizing human activities from partially observed videos. In: Computer Vision and Pattern Recognition, pp. 2658–2665 (2013)
7. Das, P., Xu, C., Doell, R.F., Corso, J.J.: A thousand frames in just a few words: Lingual description of videos through latent topics and sparse object stitching. In: Computer Vision and Pattern Recognition, pp. 2634–2641 (2013)
8. Dominey, P.F., Boucher, J.D.: Learning to talk about events from narrated video in a construction grammar framework. Artificial Intelligence **167**(1-2), 31–61 (2005)
9. Efros, A., Berg, A., Mori, G., Malik, J.: Recognizing action at a distance. In: International Conference on Computer Vision, pp. 726–733 (2003)
10. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL visual object classes

(VOC) challenge. International Journal of Computer Vision **88**(2), 303–338 (2010)
11. Everts, I., van Gemert, J.C., Gevers, T.: Evaluation of color STIPs for human action recognition. In: Computer Vision and Pattern Recognition, pp. 2850–2857 (2013)
12. Fernández Tena, C., Baiget, P., Roca, X., Gonzàlez, J.: Natural language descriptions of human behavior from video sequences. In: Advances in Artificial Intelligence, pp. 279–292 (2007)
13. Gaidon, A., Harchaoui, Z., Schmid, C.: Activity representation with motion hierarchies. International Journal of Computer Vision **107**(3), 219–238 (2014)
14. Gopalan, R.: Joint sparsity-based representation and analysis of unconstrained activities. In: Computer Vision and Pattern Recognition, pp. 2738–2745 (2013)
15. Guadarrama, S., Krishnamoorthy, N., Malkarnenkar, G., Mooney, R., Darrell, T., Saenko, K.: Youtube2text: Recognizing and describing arbitrary activities using semantic hierarchies and zero-shot recognition. In: International Conference on Computer Vision, pp. 2712–2719 (2013)
16. Gupta, A., Davis, L.S.: Objects in action: An approach for combining action understanding and object perception. In: Computer Vision and Pattern Recognition, pp. 1–8 (2007)
17. Hanckmann, P., Schutte, K., Burghouts, G.J.: Automated textual descriptions for a wide range of video events with 48 human actions. In: European Conference on Computer Vision Workshops, pp. 372–380 (2012)
18. Ikizler-Cinbis, N., Sclaroff, S.: Object, scene and actions: Combining multiple features for human action recognition. In: European Conference on Computer Vision, pp. 494–507 (2010)
19. Jain, A., Gupta, A., Rodriguez, M., Davis, L.S.: Representing videos using mid-level discriminative patches. In: Computer Vision and Pattern Recognition, pp. 2571–2578 (2013)
20. Jhuang, H., Serre, T., Wolf, L., Poggio, T.: A biologically inspired system for action recognition. In: International Conference on Computer Vision, pp. 1–8 (2007)
21. Ke, Y., Sukthankar, R., Hebert, M.: Event detection in crowded videos. In: International Conference on Computer Vision, pp. 1–8 (2007)
22. Khan, M.U.G., Gotoh, Y.: Describing video contents in natural language. In: Workshop on Innovative Hybrid Approaches to the Processing of Textual Data, pp. 27–35 (2012)
23. Khan, M.U.G., Zhang, L., Gotoh, Y.: Human focused video description. In: International Conference on Computer Vision Workshops, pp. 1480–1487 (2011)
24. Khan, M.U.G., Zhang, L., Gotoh, Y.: Towards coherent natural language description of video streams. In: International Conference on Computer Vision Workshops, pp. 664–671 (2011)
25. Kojima, A., Tamura, T., Fukunaga, K.: Natural language description of human activities from video images based on concept hierarchy of actions. International Journal of Computer Vision **50**(2), 171–184 (2002)
26. Krishnamoorthy, N., Malkarnenkar, G., Mooney, R.J., Saenko, K., Guadarrama, S.: Generating natural-language video descriptions using text-mined knowledge. In: Conference on Artificial Intelligence, pp. 541–547 (2013)
27. Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., Serre, T.: HMDB: A large video database for human motion recognition. In: International Conference on Computer Vision, pp. 2556–2563 (2011)
28. Laptev, I.: On space-time interest points. International Journal of Computer Vision **64**(2-3), 107–123 (2005)
29. Laptev, I., Marszałek, M., Schmid, C., Rozenfeld, B.: Learning realistic human actions from movies. In: Computer Vision and Pattern Recognition, pp. 1–8 (2008)
30. Le, Q.V., Zou, W.Y., Yeung, S.Y., Ng, A.Y.: Learning hierarchical invariant spatio-temporal features for action recognition with independent subspace analysis. In: Computer Vision and Pattern Recognition, pp. 3361–3368 (2011)
31. Li, L.J., Fei-Fei, L.: What, where and who? Classifying events by scene and object recognition. In: International Conference on Computer Vision, pp. 1–8 (2007)
32. Lin, Z., Jiang, Z., Davis, L.: Recognizing actions by shape-motion prototype trees. In: International Conference on Computer Vision, pp. 444–451 (2009)
33. Liu, H., Feris, R., Sun, M.T.: Benchmarking datasets for human activity recognition. In: T.B. Moeslund, A. Hilton, V. Krüger, L. Sigal (eds.) Visual Analysis of Humans: Looking at People, chap. 20, pp. 411–427. Springer (2011)
34. Liu, J., Kuipers, B., Savarese, S.: Recognizing human actions by attributes. In: Computer Vision and Pattern Recognition, pp. 3337–3344 (2011)
35. Liu, J., Luo, J., Shah, M.: Recognizing realistic actions from videos "in the wild". In: Computer Vision and Pattern Recognition, pp. 1996–2003 (2009)
36. Maji, S., Bourdev, L., Malik, J.: Action recognition from a distributed representation of pose and appearance. In: Computer Vision and Pattern Recognition, pp. 3177–3184 (2011)
37. Marszałek, M., Laptev, I., Schmid, C.: Actions in context. In: Computer Vision and Pattern Recognition, pp. 2929–2936 (2009)
38. Messing, R., Pal, C., Kautz, H.: Activity recognition using the velocity histories of tracked keypoints. In: International Conference on Computer Vision, pp. 104–111

(2009)
39. Moore, D.J., Essa, I.A., Heyes, M.H.: Exploiting human actions and object context for recognition tasks. In: International Conference on Computer Vision, pp. 80–86 (1999)
40. Niebles, J.C., Chen, C.W., Fei-Fei, L.: Modeling temporal structure of decomposable motion segments for activity classification. In: European Conference on Computer Vision, pp. 392–405 (2010)
41. Oh, S., Hoogs, A., Perera, A., Cuntoor, N., Chen, C.C., Lee, J.T., Mukherjee, S., Aggarwal, J.K., Lee, H., Davis, L., Swears, E., Wang, X., Ji, Q., Reddy, K., Shah, M., Vondrick, C., Pirsiavash, H., Ramanan, D., Yuen, J., Torralba, A., Song, B., Fong, A., Roy-Chowdhury, A., Desai, M.: A large-scale benchmark dataset for event recognition in surveillance video. In: Computer Vision and Pattern Recognition, pp. 3153–3160 (2011)
42. Oneata, D., Verbeek, J., Schmid, C.: Action and event recognition with Fisher vectors on a compact feature set. In: International Conference on Computer Vision, pp. 1817–1824 (2013)
43. Oreifej, O., Shah, M.: Robust Subspace Estimation Using Low-Rank Optimization, chap. 5, pp. 55–67. Springer (2014)
44. Reddy, K.K., Shah, M.: Recognizing 50 human action categories of web videos. Machine Vision and Applications **24**(5), 971–981 (2013)
45. Rodriguez, M.D., Ahmed, J., Shah, M.: Action MACH: A spatio-temporal maximum average correlation height filter for action recognition. In: Computer Vision and Pattern Recognition, pp. 1–8 (2008)
46. Rohrbach, M., Qiu, W., Titov, I., Thater, S., Pinkal, M., Schiele, B.: Translating video content to natural language descriptions. In: International Conference on Computer Vision, pp. 433–440 (2013)
47. Ryoo, M.S.: Human activity prediction: Early recognition of ongoing activities from streaming videos. In: International Conference on Computer Vision, pp. 1036–1043 (2011)
48. Sadanand, S., Corso, J.J.: Action bank: A high-level representation of activity in video. In: Computer Vision and Pattern Recognition, pp. 1234–1241 (2012)
49. Schuldt, C., Laptev, I., Caputo, B.: Recognizing human actions: A local SVM approach. In: International Conference on Pattern Recognition, pp. 32–36 (2004)
50. Siddharth, N., Barbu, A., Siskind, J.M.: Seeing what you're told: Sentence-guided activity recognition in video. In: Computer Vision and Pattern Recognition, pp. 732–739 (2014)
51. Siskind, J.M.: Visual event classification via force dynamics. In: Conference on Artificial Intelligence, pp. 149–155 (2000)
52. Siskind, J.M., Morris, Q.: A maximum-likelihood approach to visual event classification. In: European Conference on Computer Vision, pp. 347–360 (1996)
53. Smeaton, A., Over, P., Kraaij, W.: Evaluation campaigns and TRECVID. In: ACM International Conference on Multimedia Information Retrieval, pp. 321–330 (2006)
54. Song, Y., Morency, L.P., Davis, R.: Action recognition by hierarchical sequence summarization. In: Computer Vision and Pattern Recognition, pp. 3562–3569 (2013)
55. Soomro, K., Zamir, A.R., Shah, M.: UCF101: A dataset of 101 human actions classes from videos in the wild. Tech. Rep. 1212.0402, arXiv (2012)
56. Tang, K., Fei-Fei, L., Koller, D.: Learning latent temporal structure for complex event detection. In: Computer Vision and Pattern Recognition, pp. 1250–1257 (2012)
57. Tian, Y., Sukthankar, R., Shah, M.: Spatiotemporal deformable part models for action detection. In: Computer Vision and Pattern Recognition, pp. 2642–2649 (2013)
58. Uemura, H., Ishikawa, S., Mikolajczyk, K.: Feature tracking and motion compensation for action recognition. In: British Machine Vision Conference, pp. 1–10 (2008)
59. Wang, C., Wang, Y., Yuille, A.L.: An approach to pose-based action recognition. In: Computer Vision and Pattern Recognition, pp. 915–922 (2013)
60. Wang, H., Kläser, A., Schmid, C., Liu, C.L.: Action recognition by dense trajectories. In: Computer Vision and Pattern Recognition, pp. 3169–3176 (2011)
61. Wang, H., Kläser, A., Schmid, C., Liu, C.L.: Dense trajectories and motion boundary descriptors for action recognition. International Journal of Computer Vision **103**(1), 60–79 (2013)
62. Wang, H., Schmid, C.: Action recognition with improved trajectories. In: International Conference on Computer Vision, pp. 3551–3558 (2013)
63. Wang, Y., Mori, G.: Human action recognition by semilatent topic models. IEEE Transactions on Pattern Analysis and Machine Intelligence **31**(10), 1762–1674 (2009)
64. Wang, Z., Kuruoglu, E.E., Yang, X., Xu, Y., Yu, S.: Event recognition with time varying hidden Markov model. In: International Conference on Accoustic and Speech Signal Processing, pp. 1761–1764 (2009)
65. Willems, G., Tuytelaars, T., Van Gool, L.: An efficient dense and scale-invariant spatio-temporal interest point detector. In: European Conference on Computer Vision, pp. 650–663 (2008)
66. Wu, X., Xu, D., Duan, L., Luo, J.: Action recognition using context and appearance distribution features. In: Computer Vision and Pattern Recognition, pp. 489–496 (2011)

67. Xu, G., Ma, Y.F., Zhang, H., Yang, S.: Motion based event recognition using HMM. In: International Conference on Pattern Recognition, pp. 831–834 (2002)
68. Yamoto, J., Ohya, J., Ishii, K.: Recognizing human action in time-sequential images using hidden Markov model. In: Computer Vision and Pattern Recognition, pp. 379–385 (1992)
69. Yu, H., Siskind, J.M.: Grounded language learning from video described with sentences. In: Annual Meeting of the Association for Computational Linguistics, pp. 53–63 (2013)
70. Yuan, C., Hu, W., Tian, G., Yang, S., Wang, H.: Multi-task sparse learning with Beta process prior for action recognition. In: Computer Vision and Pattern Recognition, pp. 423–429 (2013)
71. Yuan, C., Li, X., Hu, W., Ling, H., Maybank, S.: 3D R transform on spatio-temporal interest points for action recognition. In: Computer Vision and Pattern Recognition, pp. 724–730 (2013)
72. Yuan, J., Liu, Z., Wu, Y.: Discriminative subvolume search for efficient action detection. In: Computer Vision and Pattern Recognition, pp. 2442–2449 (2009)
73. Zhu, J., Wang, B., Yang, X., Zhang, W., Tu, Z.: Action recognition with Actons. In: International Conference on Computer Vision, pp. 3559–3566 (2013)